Micro-burst in Data Centers: Observations, Implications, and Applications

Danfeng Shan*, Fengyuan Ren*, Peng Cheng*†, Ran Shu*
* Tsinghua University † Microsoft Research Asia

arXiv:1604.07621v1 [cs.NI] 26 Apr 2016

ABSTRACT

Micro-burst traffic is not uncommon in data centers. It can cause packet dropping, which results in serious performance degradation (e.g., Incast problem). However, current solutions that attempt to suppress micro-burst traffic are extrinsic and ad hoc, since they lack the comprehensive and essential understanding of micro-burst's root cause and dynamic behavior. On the other hand, traditional studies focus on traffic burstiness in a single flow, while in data centers micro-burst traffic could occur with highly fan-in communication pattern, and its dynamic behavior is still unclear.

To this end, in this paper we re-examine the micro-burst traffic in typical data center scenarios. We find that evolution of micro-burst is determined by both TCP's self-clocking mechanism and bottleneck link. Besides, dynamic behaviors of micro-burst under various scenarios can all be described by the slope of queue length increasing. Our observations also implicate that conventional solutions like absorbing and pacing are ineffective to mitigate micro-burst traffic. Instead, senders need to slow down as soon as possible. Inspired by the findings and insights from experimental observations, we propose S-ECN policy, which is an ECN marking policy leveraging the slope of queue length increasing. Transport protocols utilizing S-ECN policy can suppress the sharp queue length increment by over 50%, and reduce the 99^{th} percentile of query completion time by $\sim$20%.

Keywords

Micro-burst traffic; Packet dropping; TCP; Queue length; Switch buffer

1. INTRODUCTION

Broadly speaking, micro-burst is highly intense traffic appearing in a relatively short period. Recently, micro-bursts in modern data centers attract much attention in both academic and industry community since they cause serious performance problems. When massive packets swarm into the same switch port in a twinkling of an eye, the buffer is sharply bulged, and even overwhelmed. The resulting high queuing delays and jitters impose heavy penalties on financial trading applications in data centers [3], and the packet dropping in succession triggers abnormal timeouts, which causes TCP incast throughput collapse [35, 42, 14], and sluggish response time or low quality of results [6, 44, 41].

Because of a lack of comprehensive and essential understanding on both root cause and dynamic behavior of micro-burst in data centers, most of existed solutions to suppress micro-burst and eliminate its negative impact are ad hoc and extrinsic, for example reducing RTOmin [42], limiting the number of concurrent flows and adding jitter in application layer [37], absorbing micro-burst in switch buffer [9, 19, 21, 39], smoothing micro-burst through pacing packets at sources [7].

Intuitively, micro-burst traffic stems from various elements, including data generation pattern in applications [26, 30], system call batching and protocol stack processing in OS [30], coalescing packets in NICs [32, 16]. Definitely, the end-to-end congestion control mechanism built-in window-based transport protocol, such as TCP and DCTCP, is the source of a considerable amount of traffic bursts. Historically, lots of work [50, 17, 26, 27, 13, 8, 38] examined the nature of both micro-burst and macro-burst [8] induced by transport protocols in the context of traditional Internet, and mainly focused on the effect of a single long-lived TCP flow on the formation and dynamic evolution of burst traffic.

The features of modern data center networks are obviously different from the counterparts of traditional Internet, such as symmetrical topology and short network radius. The traffic distribution and pattern are also unique, for example, the short-lived flows are numerous, but a few long-lived flows monopolize the majority of network resources [28]. The new computing paradigms (such as MapReduce [20] and distributed streaming computing [33]) and application systems (such as in-memory computing [34], and distributed machine learning [31]) require various communication patterns, which intensify high fan-in bursty traffic [9]. Moreover, diverse workflow patterns (such as Partition/Aggregate and Dependent/Sequential [48, 29]) further complicate traffic patterns in data centers. Due to the above factors, the conclusions, which describe the nature of micro-burst and macro-burst induced by a single long TCP

flow in Internet, cannot be extended to predicting the dynamic behavior of micro-burst caused by aggregated TCP short flows in data centers, and thus hardly inspire some reasonable guidelines towards removing its negative effects radically.

Considering the features of traffic pattern and distribution in data centers, in this work, we re-examine the micro-burst traffic under the control of window-based transport protocol. The main contributions are summarized in three aspects:

(1) Monitoring the evolution of queue length at fine-grained timescales in typical traffic scenarios and conducting analysis on observations, we infer the temporal and spatial dynamic behaviors of micro-burst, and obtain some interesting findings: 1) The self-clocking mechanism and the bottleneck link jointly dominate the evolution of micro-burst; 2) In any traffic scenarios, regardless of with or without long-lived background flows, and including synchronous or asynchronous fan-in short-lived flows, there is an immutable variable (i.e., the slope of queue length increasing) that can precisely describe the dynamic behaviors of micro-burst in most situations.

(2) Our findings indicate some implications about guidelines towards suppressing micro-burst. The conventional mechanisms, such as absorbing and pacing, are ineffective in the context of data centers. Radically, as for latency-sensitive short message transmission in data centers, the ideal solution is to detect congestion in time and to properly throttle sending rate as soon as possible.

(3) Enlightened by our findings and implications, and leveraging the slope of queue length increasing as an indicator of detecting congestion, we propose an ECN probability marking policy, called S-ECN, to eliminate the negative impact of micro-burst caused by aggregated concurrent flows. Since the slope of queue length can predict the dynamic behavior of micro-burst, when the S-ECN policy is employed, the queue length increment is effectively suppressed by over 50% compared with mainstream transport protocols in data centers. Overall, 99^{th} percentile of query completion time can be reduced by ∼20%.

The rest of paper is organized as follows. In §2, we briefly describe the background and present our methodology. In §3, we study the micro-burst traffic under different scenarios, and show our experimental observations and findings. In §4, we propose S-ECN policy and evaluate its performance. Finally, we discuss some related work in §5, and conclude in §6.

2. METHODOLOGY

In this section, we discuss the scenarios we need to consider, and describe how we observe micro-burst traffic.

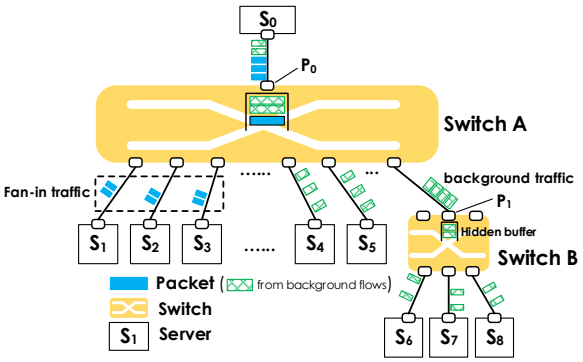

Figure 1: Micro-burst in data centers

2.1 Scenario Classification

As depicted in Figure 1, to comprehensively study the nature and evolution mechanism of micro-burst traffic under the control of transport protocol in data centers, five typical scenarios are summarized as follows.

(1) *Synchronous fan-in traffic:* In this scenario, many concurrent flows start simultaneously and destine for the same target, as depicted in Figure 1. This scenario is common in many data center application systems. For example, in distributed storage system [35], a client will synchronously read data blocks from lots of storage servers. In user-facing online services employing partition-aggregate structure [6], a parent server will query results from many leaf servers in parallel, and leaf servers will synchronously sending responses to the parent server. In memcached systems [34], a web server will communicate with many memcached servers to satisfy a user request.

(2) *Asynchronous fan-in traffic:* Different from previous scenario, in this scenario, many concurrent flows start asynchronously, but destine for the same target. This scenario also exists in multiple data center applications. For example, in distributed machine learning systems [31], lots of servers frequently communicate with a centralized server, and thus communications from different sources easily overlap although they are asynchronous. And in MapReduce framework [20], many mappers will transfer intermediate results to reducers, and these transmissions may be asynchronous but can still be bottlenecked together.

(3) *Fan-in traffic with one background flow:* In this scenario, there is one long-lived background flow before fan-in flows start. The background flow will share the same bottleneck with fan-in flows (congested at port P_0 in Figure 1).

(4) *Fan-in traffic with several background flows congested at the same hop:* As depicted in Figure 1, in this scenario, several background flows have already congested at the port P_0 before fan-in flows start. Different from previous scenario, when fan-in flows start, they

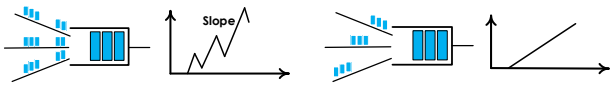

(a) Synchronous arrival (b) Asynchronous arrival

Figure 2: Characteristics of micro-burst traffic can be inferred from queue length evolution.

can see buffer occupancy caused by background flows at port P_0, which may affect the dynamic behavior of micro-burst traffic.

(5) *Fan-in traffic with several background flows congested at previous hop:* As depicted in Figure 1, in this scenario, several background flows are congested at port P_1 in the beginning. When fan-in flows arrive at port P_0, they cannot see any buffer occupancy. However, as the congestion point shifts from port P_1 in Switch B to port P_0 in Switch A, the buffer occupancy in the port P_1 will move to the port P_0 in a very short time, which may intensify the traffic burstiness at port P_0.

2.2 Observing Micro-burst

Capturing or monitoring micro-bursts in data centers is challenging [18, 40], since switches operate at very high speed (10/40Gbps) while micro-burst occurs in very short timescales (hundreds of microseconds). On the other hand, through separately studying individual micro-burst on each link or at end systems, we cannot directly gain the dynamic behavior of micro-burst when lots of flows aggregate, which is significant since micro-burst traffic causes problems mainly when fan-in traffic causes fast increasing of queue length.

In this paper, we study micro-burst traffic through observing evolution of queue length in switch, which directly reflects the aggregation behavior of traffic. Furthermore, since queue length in switch evolves differently under different scenarios, the dynamic behavior of micro-burst caused by aggregation of flows can be inferred from queue length evolution, as shown in Figure 2. For example, when fan-in flows start synchronously, bursts from different flows may arrive simultaneously. As a result, there are impulses during queue length increasing (Figure 2(a)). When fan-in flows start asynchronously, bursts from different flows may arrive separately. As a result, queue length increases in a smoother way (Figure 2(b)). Besides, queue length evolution is also different when there are background flows. Moreover, as depicted in Figure 2(a), slope of queue increasing (i.e., queue length increasing rate) can reflect the fan-in degree, since with more concurrent flows the queue length tends to increase faster.

2.3 Monitoring Queue Length

To observe micro-burst traffic in switch queue, we need to get queue length at fine-grained timescales, because the queue length changes every a few microseconds. Existing techniques in commodity switches can only monitor queue length at a coarse timescale. Thus, we develop a method similar to TPP [25], which can acquire queue length at per-packet granularity, and introduce little overhead to switch in the meanwhile. Specifically, whenever a packet is passing through the switch, we let it carry the timestamp and current queue length in its payload.[1] When the packet arrives at the receiver, the receiver extracts timestamp and queue length from it, and stores the information into disk.

We implement the approach on NetFPGA platform [4]. Specifically, we use a 16-bit register to record the queue length in each port, which is at an 8-byte granularity. We also use a 32-bit register to maintain timestamp, which is at an 800-ns granularity (i.e., the register increases every 800ns). Timestamp and queue length are put into packets before they enter the output queue, which only need additional processing time of 1 cycle. Since the clock run at 125MHz, this implementation only adds 8ns latency to the pipeline. Our implementation needs less than 6% extra resources, as shown in Table 1.

Resources	Ref. Switch	ECN Switch	+TPP	+S-ECN
Slice Flip Flops	14158	14378	14777	14700
LUTs	17589	18048	19050	18544

Table 1: Resource usage in NetFPGA. ECN switch is built upon reference switch. TPP switch and S-ECN switch are built upon ECN switch.

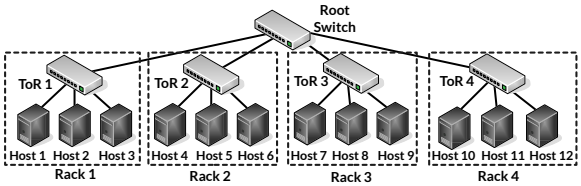

Figure 3: Network topology in experiments

3. EXPERIMENTAL OBSERVATIONS

In this section, we separately study micro-burst traffic in each scenario listed in §2.1. After that, we present our observations and implications.

3.1 Testbed

The topology of our testbed is shown in Figure 3. The testbed contains 12 hosts across 4 racks. Each rack

[1]In order not to affect normal traffic (such as ssh) that carries useful data in the payload, only packets, whose port value (both source port and destination port) in TCP header is larger than 1023, are used to carry queue length.

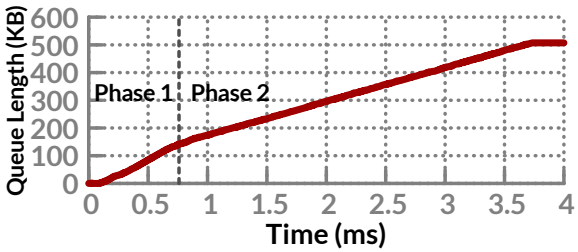

(a) Queue length evolution

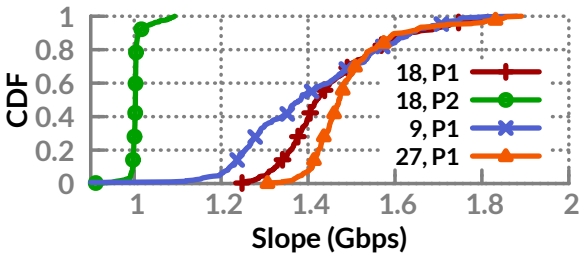

(b) Distribution of slope

Figure 4: Experiment results in synchronous fan-in scenario. In (b), "18,P1" stands for slope in Phase 1 when there are 18 flows.

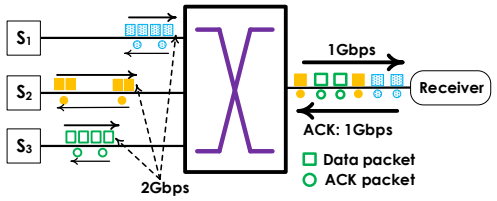

Figure 5: Explanation of Phase 2

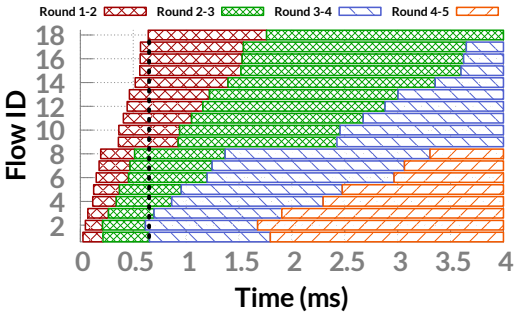

Figure 6: Time interval between two successive rounds of packets

holds 3 hosts connected to a top-of-rack (ToR) switch. All ToR switches are connected through a root switch. Each host is a Dell Optiplex 780 desktops with an Intel® Core™ 2 Duo E7500 2930 MHz CPU, 4 GB memory, a 500GB hard disk, and an Intel® 82567LM Gigabit Ethernet NIC, running CentOS 5.11 with GNU/Linux kernel 2.6.38. All switches are NetFPGA cards with four Gigabit Ethernet networking ports. Buffer size in each output port can be arbitrarily set between 1KB and 512KB. RTT between hosts from different racks is about 50μs without queuing. The queue length in root switch is monitored.

3.2 Observations and Analysis

3.2.1 *Synchronous fan-in traffic*

Experiment settings: Firstly, we let Host 10 send queries to Host 1-9 at the same time and these hosts will simultaneously send a respond message to Host 10. Each host is used to emulate multiple senders. Query size is 50B and response message size is 1000KB. The buffer size in each output port is set to 512KB. End hosts use TCP NewReno as their transport protocols[2]. Large Segment Offload is disabled, but it will not affect our main results (see discussion in §3.5). Delayed ACK is also disabled [47]. We use default values for other TCP parameters.

Experiment results: Figure 4(a) depicts the queue length evolution in the congested port when there are 18 senders. The queue length increasing has two phases: sharply increasing for a relatively short period at the beginning (Phase 1) and slower increasing after that (Phase 2). Particularly, in Phase 2, the queue length increases at a constant rate, without any impulses or jitters.

We repeat the experiment for 100 times and calculate the slope in Phase 1 and Phase 2, respectively. The slope distributions are shown in Figure 4(b), where we have another observation: in Phase 2, *the queue length is not only increasing at a constant rate but the increasing rate is always 1Gbps.* Specifically, when there are 18 flows, the average slope is 1.005Gbps, and the coefficient of variation is only 0.018. Through detailed analysis, we find that this phenomenon is caused by TCP's self-clocking system and evolution of congestion window in slow start phase, as depicted in Figure 5. Specifically, since data packets are received at 1Gbps, the receiver is always acknowledging data at 1Gbps, which is the speed of congested port. On the other hand, since all flows are in slow start phase, each sender introduces two data packets into network when receiving an ACK packet (note that delayed ACK is disabled). As a result, the total generated traffic from all senders is 2Gbps. Thus, the overall queue length increasing rate is 1Gbps. Furthermore, at packet granularity, ACK packets are evenly spread across bottleneck link. Since the topology is symmetric, after these ACKs arrive at senders, they will evenly trigger transmissions of data packets, and thus data packets are evenly spread across input links. As a result, after these packets aggregate, the queue length in congested port is smoothly increasing

[2]We have studied different TCP variants (Reno, SACK, CUBIC) and observed the similar results.

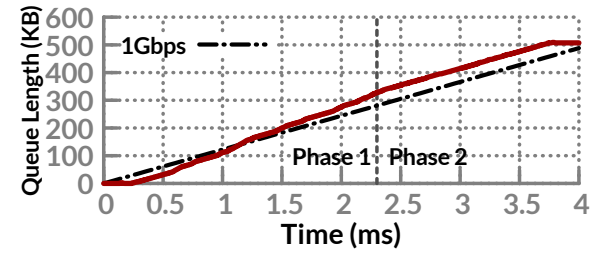

(a) Queue length evolution

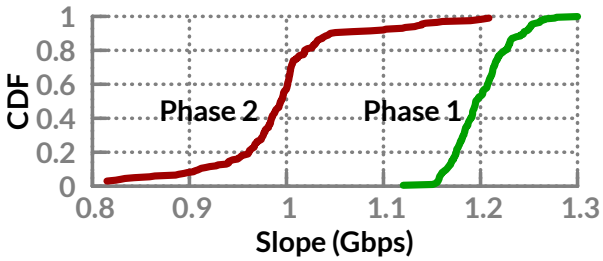

(b) Distribution of slope

Figure 7: Experiment results in asynchronous fan-in scenario.

at constant rate.

As shown in Figure 4(b), in Phase 1, the slope is much larger than that in Phase 2; and with more concurrent flows, the slope in Phase 1 tends to be larger. This is because sources are sending the first round of packets in this phase, and the self-clocking system does not take effect. To prove this, we find out the ending time of first round among all flows. Figure 6 shows the time interval between ending time of two successive rounds. Specifically, the latest ending time of transmitting first round packets is at 0.642ms, which is in accordance with the ending time of Phase 1 in Figure 4(a).

3.2.2 *Asynchronous fan-in traffic*

Experiment settings: In this experiment, 18 flows randomly start during a period of 2ms, which are from Host 1-9 to Host 10. Other settings keep unchanged.

Experiment results: The queue length evolution is shown in Figure 7(a). Because first rounds of packets from different senders asynchronously arrive at the congested port, in Phase 1 the queue length does not increase as sharply as that in previous scenario where these packets arrive synchronously. However, in Phase 2, the queue length is smoothly increasing at a constant rate of port speed, which is similar to that in synchronous. This is because all flows have finished sending the first round of packets, and all packet departures in this phase are triggered by ACKs. In addition, as in previous scenario, ACKs are evenly acknowledging data at 1Gbps, and each ACK packet triggers transmissions of two packets since flows are in slow start phase. Therefore, traffic arriving rate at the congested port is 2Gbps and slope is 1Gbps.

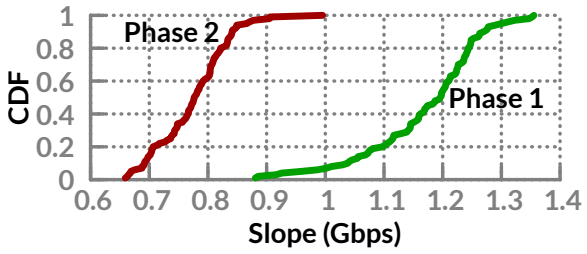

Figure 8: Distribution of slope when there is one background flow.

3.2.3 *Fan-in traffic with one background flow*

Experiment settings: Along with query traffic, we add one background flow. Specifically, at the beginning, Host 9 begins to send data to Host 11 as fast as it can. After 0.5 seconds, Host 10 begins to query Host 1-8. Other settings are not changed.

Experiment results: The experiment is repeated for 100 times. As shown in Figure 8, the slope in Phase 1 is larger than the port speed as in previous scenarios, but the slope in Phase 2 becomes much smaller than the port speed. This is because the congestion window of the background flow has reached its maximum value by the time other fan-in flows start, and thus the background flow only introduces one data packet into network when receiving an ACK packet[3]. On the other hand, the other flows inject two data packets when receiving an ACK. Overall, the arriving rate of total traffic is less than 2Gbps, and thus slope is smaller than 1Gbps.

3.2.4 *Fan-in traffic with several background flows congested at the same hop*

Experiment settings: Along with query traffic, we add three background flows, which are from Host 1, 4, 7 to Host 11 and Host 12. After 0.5 seconds, Host 10 begins to query hosts in other racks (excluding Host 1, 4, 7) and these hosts send a response message to Host 10.

Experiment results: The queue length evolution is shown in Figure 9(a). The queue length increasing can be split into multiple phases. Before the fan-in flows start, there is buffer occupancy caused by background flows in the root switch. In Phase 1, first round of packets from the fan-in flows arrive at the port simultaneously, therefore the queue length experiences a sharp in-

[3] Actually, at the beginning, the sender of the background flow keeps sending data at line rate. Because the maximum congestion window (which is usually limited by TCP's send buffer size) can be larger than the pipeline capacity ($C \times RTT$). But this has little effect as long as TCP's send buffer is not too large. Therefore, we don't discuss it in this paper.

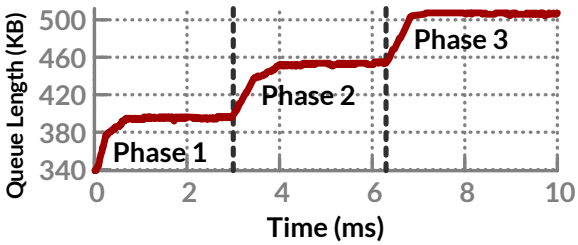

(a) Queue length evolution.

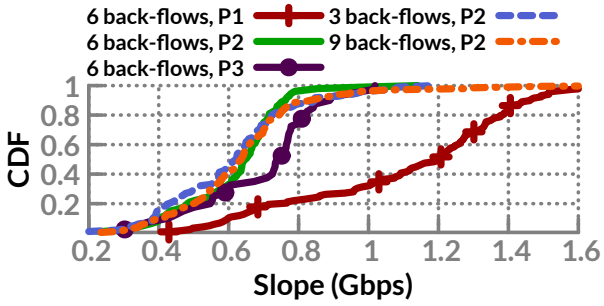

(b) Distribution of slope. "P2" stands for "Phase 2".

Figure 9: Experiment results when there are several background flows congested at the same hop.

creasing. After that, since the packets from fan-in flows are queued behind those from background flows, fan-in flows will defer injection of new packets into network, and the queue length stops sharply increasing. Similarly, in Phase 2, the queue length increases sharply at the beginning, which is caused by the second round of packets, and the queue length keeps smoothly increasing after that. Finally, fast queue length increasing in Phase 3 is caused by the third round of packets.

We repeat the experiment for 100 times, and calculate the slope in each phase when queue length is sharply increasing, which is shown in Figure 9(b). We have three main observations. *First,* we find that slope after Phase 1 (i.e., Phase 2 and Phase 3) is much smaller than that without background flows. This is because some of the arriving packets in these phases are from background flows, and these background flows are in congestion avoidance phase, in which the congestion window is increased by one MSS every round trip time. Therefore, although ACKs are acknowledging data at 1Gbps, data packets are sent at a rate smaller than 2Gbps. *Second,* after Phase 1, the slope in later phase (e.g., Phase 3) is larger than that in earlier phase (e.g., Phase 2). This is because ACKs from slow start flows introduce more packets into network than those from flows in congestion avoidance phase, and the ratio of packets from slow start flows becomes increasingly large. As a result, in the reverse path, the ratio of ACKs from slow start flows also becomes increasingly large.

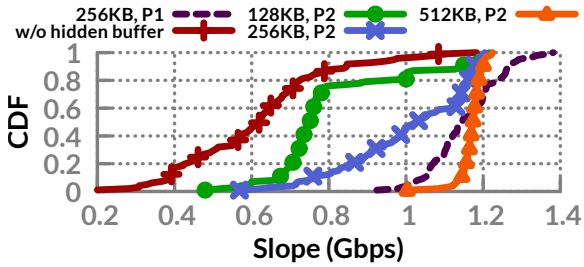

Figure 10: Slope distribution when there is hidden buffer. "P2" stands for "Phase 2".

Finally, after these ACKs arrive at sender, the total packet departure rate is larger in later phase. *Third,* the slope distributions are similar with different number of background flows. For example, the average slopes are (0.613Gbps, 0.622Gbps, 0.641Gbps) when there are (3, 6, 9) background flows. This is because the ratio of packets from slow start phase has little change as the number of background flows varies, and the congestion window increasing rate from background flows can be ignored compared to that from slow start flows.

3.2.5 *Fan-in traffic with several background flows congested at previous hop*

Experiment settings: In this experiment, the background flows are from Host 7, 8, 9 to Host 11 and Host 12, thus, these background flows are congested in ToR switch 3. After 0.5 seconds, Host 10 begins to query Host 1-6 and they send a respond message to Host 10. At this time, the congestion point will shift to the next hop in root switch. Other experiment settings keep unchanged.

Experiment results: We repeat each experiment for 100 times. Figure 10 shows the distribution of slope. We compare the slope distribution with that in previous experiment in which the congestion point does not change. Experiment results show that in this scenario slope in Phase 2 is much larger than that without shift of congestion point, and may even be larger than the port speed. This phenomenon is caused by the buffer occupancy in previous hop. Specifically, in Phase 2, arriving packets at the root switch are from fan-in flows and background flows. The packet transmissions from fan-in flows follow self-clocking mechanism. However, the transmissions of packets from background flows are not determined by self-clocking mechanism, since they are from buildup queue in the previous hop (i.e., ToR switch 3), and they are sent at the port speed of ToR switch 3. As a result, the queue length increasing is intensified by these packets. Since the buffer occupancy in previous hop cannot be seen by fan-in flows at the beginning, we call it *hidden buffer*.

We examine the effect of hidden buffer by setting the buffer size of ToR switch 3 to 128KB, 256KB, and

512KB, respectively. As shown in Figure 10, the slope increases as the amount of packets in hidden buffer increases. Specifically, when buffer sizes are 128KB, 256KB, and 512KB, the average slopes are 0.804Gbps, 1.009Gbps, and 1.170Gbps, respectively.

In theory, the slope is no larger than $2aR/(a+R)$, where a is the arriving rate of first round packets from fan-in flows and R is the port speed. This is because after 1st round packets arrive at congested port, the fraction of bandwidth occupied by fan-in flows is about $aR/(a+R)$. Therefore, in the next round, these fan-in flows will generate traffic at $2aR/(a+R)$ in total. Meanwhile, sending rate of traffic from hidden buffer to the congested port is at most R. Therefore, the queue length increasing rate is no larger than $2aR/(a+R)$.

Note that although in our experiments the packets in hidden buffer only come from one input port, results can be extended to the scenario in which the packets in hidden buffer come from several input ports, because total arriving rate of traffic from hidden buffer is no larger than the port speed. Otherwise, the bottleneck is not in the previous hop. In particular, the upper bound of slope is also $2aR/(a+R)$.

3.3 Summary

In summary, when multiple flows aggregate in a queue, the queue length increasing can be divided into two phases. In Phase 1, queue length increasing is caused by arriving of 1st round packets. The slope in Phase 1 directly reflects the communication pattern and intensity. Specifically, when communications are synchronous, or more concurrent flows aggregate in a queue, the slope in this phase should be larger.

In Phase 2, we have two main observations. From a macroscopic perspective, the arriving traffic rate is limited by the bottleneck speed, because total amount of sent data among all bottlenecked flows is always acknowledged at bottleneck rate. Besides, the traffic arriving rate is directly reflected by the slope of queue length increasing. Under different traffic scenarios, dynamics of slope can be summarized as following laws:

Law 1 Without background flows, the slope in Phase 2 is equal to the port speed.

Law 2 When there is one background flow, or several background flows are congested at the same hop, the slope in Phase 2 is smaller than the port speed.

Law 3 When several background flows are congested at previous hop, slope in Phase 2 might be larger than port speed, but no larger than $2aR/(a+R)$, where a is the arriving rate of 1st round packets and R is the port speed.

Moreover, from a microscopic perspective, packet arrival pattern in forward path is evenly scheduled by ACK packets in reverse path. Specifically, data packets from different sources are sent at full rate to bottleneck link. After they arrive at the receiver, the ACKs will evenly spread in the reverse path of bottleneck link. After these ACKs are fanned out to multiple senders, they will evenly trigger the transmissions of data packets into network. As a result, in the bottleneck queue, the queue length increases smoothly.

In conclusion, our experimental observations reveal that the bottleneck link and the self-clocking system are playing an important role in the evolution of micro-burst traffic.

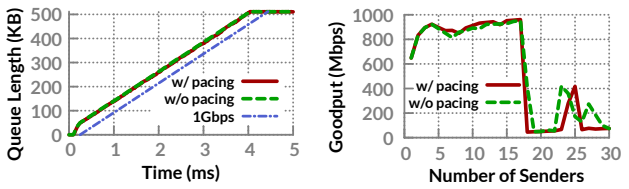

(a) Queue length evolution (b) Incast performance

Figure 11: Performance of TCP pacing. In the simulation, 40 servers are connected through the same switch. Experiment workload: (a) The same as that in §3.2.1. (b) Each sender sends 64KB data to the same receiver. Switch buffer size is 128KB.

3.4 Implications

These experimental observations have some implications:

(1) Slope is an important indicator to mitigate micro-burst traffic. Slope directly reflects the bursty degree of aggregated traffic. Specifically, the slope in Phase 1 reflects the flow concurrency, and the slope in Phase 2 reflects the mismatch between congestion window increasing rate and bottleneck rate. By carefully slowing down senders according to the slope, the sharp queue length increasing caused by micro-burst traffic can be effectively suppressed.

(2) TCP pacing does not help ease the sharp queue length increasing caused by fan-in traffic. In the literature [50, 5, 43, 27], TCP pacing is considered as an effective method to smooth traffic burstiness in a single flow by evenly spreading a window of packets over a round-trip time. However, in data centers with fan-in traffic, traffic burstiness comes from multiple flows. Besides, transmissions of data packets have already been evenly triggered by ACK packets (as is analyzed in §3.2.1). Thus TCP pacing has little effect on queue length evolution, and may not alleviate the impairments caused by micro-burst traffic (such as TCP incast throughput collapse). To validate this, we conduct simulations on ns-2 platform [2, 1]. The results are shown in Figure 11.

(3) Simply absorbing micro-burst traffic may

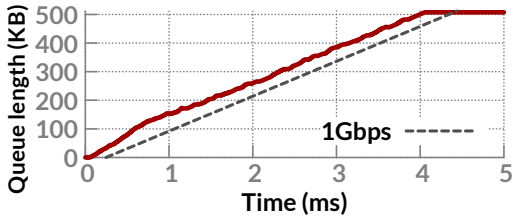

Figure 12: Queue length evolution when Large Segment Offload is turned on.

not help avoiding packet dropping, but partly magnifies micro-burst traffic. Senders won't slow down before they receive congestion signals. If packets are absorbed without notifying senders of congestion, the senders will believe that network is still fine and introduce more packets into network. Specifically, if the sender is in slow start phase, absorbing every packet will introduce another two packets into network.

3.5 Discussion

Impact of other micro-burst traffic. One might ask whether other traffic burstiness (e.g., caused by LSO) affects our observations. We find that these kinds of micro-burst traffic can only cause queue jitter at very small timescales, because they disrupt the even transmission of data packets. For example, we turn on LSO, and use the same experiment settings as §3.2.1, As shown in Figure 12, the queue length evolution is similar to that without LSO.

Is these observations specific to TCP? Our most phenomena can also be observed in other TCP-like protocols, because they use the same self-clocking mechanism and similar congestion window adjustment algorithm as TCP. Besides, the hidden buffer is caused by persistent buffer occupancy from long flows. Since most protocols can not promise zero buffer occupancy, hidden buffer also exists in other protocols. However, current data center protocols, such as DCTCP [6] and ECN* [46], usually keep buffer occupancy low, the effect of hidden buffer might be not as serious as observed in our experiments.

4. MITIGATING MICRO-BURST

From previous experiments, we find that to mitigate micro-burst traffic, senders need to slow down in time. In this section, we discuss the existing solutions that can potentially mitigate micro-burst traffic, and show their limitations. After that, we propose S-ECN scheme, which can effectively suppress sharp queue length increasing caused by micro-burst traffic, and reduce hidden buffer by keeping low buffer occupancy caused by long-lived flows.

4.1 ECN Marking

ECN [36] can explicitly notify senders of congestion before buffer overflows. Therefore, by using ECN, packet dropping caused by micro-burst traffic may be avoided. Currently, ECN is widely used by Active Queue Management (AQM) schemes (such as RED [22] and PI [24]) and some specific protocols (such as DCTCP [6] and ECN* [46]). In these schemes, packets are marked according to queue length. For example, RED marks packets according to the average queue length. DCTCP and ECN* mark packets according to instant queue length.

However, queue-length-based ECN marking has limitations when micro-burst traffic causes sharp queue length increasing, because packets are not marked until the queue length reaches a threshold. After queue length reaches the threshold, it takes time for all senders to receive the feedbacks from switch queue. During this time, the queue length still keeps sharply increasing and buffer may overflow before all senders slow down. For example, assume that packets are marked with ECN when the instant queue length is larger than the threshold. Multiple flows start concurrently and they share the same bottleneck. At the beginning, all flows are in slow start phase, therefore, the queue length will be increasing at port's sending rate (denoted by R). When queue length reaches the ECN threshold (denoted by h), the 1st sender will decrease its congestion window after $h/R + RTT$. During this period, the queue length will increase to $h+R\cdot(h/R+RTT) = 2h+R\cdot RTT$. Note at this time only the first sender slows down, and it takes much more time for all senders to slow down. We conduct an experiment when there are 9 senders (more details in §4.4). Figure 14(b) shows that the queue length does not stop increasing until it reaches 104KB with DCTCP protocol, while ECN threshold is only 32KB. On the other hand, the ECN threshold should not be too low, otherwise network cannot be fully utilized [46]. Therefore, we need another scheme to rapidly suppress the sharp queue length increasing.

4.2 Design of S-ECN

To effectively suppress the micro-burst traffic, we propose S-ECN. S-ECN is inspired by the implication of slope: slope reflects the mismatch between the congestion window increasing rate and bottleneck rate, thus sharp queue increasing can be suppressed by properly marking packets according to slope. Specifically, we mark packets immediately when the queue length begins to increase. Furthermore, the fraction of marked packets indicates how large the slope is: the bigger the slope, the larger the fraction.

Specifically, when slope (denoted by s) is equal to or even larger than the port speed (denoted by R, which is constant in general), then flows are in slow start phase. All packets are marked with ECN so that flows will slow

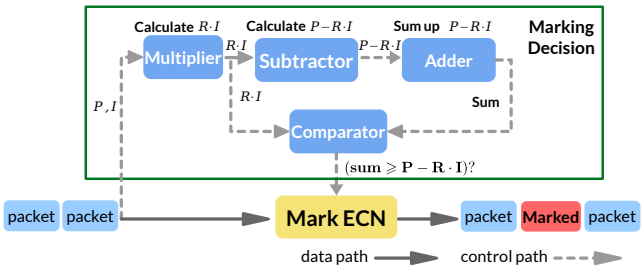

Figure 13: Implementation structure of S-ECN

Protocol Name	End Host Algorithm	Switch	
		S-ECN	ECN
TCP	TCP	OFF	OFF
ECN*	TCP	OFF	ON
S-ECN	TCP	ON	OFF
SL-ECN	TCP	ON	ON
DCTCP	DCTCP	OFF	ON
DCTCP+SL-ECN	DCTCP	ON	ON

Table 2: Protocols used in evaluation

down immediately. When slope is lower than the port's speed, then packets are marked at a probability of s/R. Finally, when slope is lower than or equal to 0 (i.e., the queue length is not increasing), none of packets are marked. The marking probability (denoted by $Prob$) as a function of slope s is as equation (1).

$$Prob = \begin{cases} 0, & s \leqslant 0, \\ \frac{s}{R}, & 0 < s < R, \\ 1, & s \geqslant R \end{cases} \quad (1)$$

4.3 Implementation of S-ECN

We implement our scheme on NetFPGA platform [4]. For simplicity, per-packet processing is used. Specifically, in NetFPGA, we can easily get following information: (i). Total size of an arriving packet (denoted by P, in bytes). (ii). Time interval between arrivals of previous packet and the current packet (denoted by I, in clocks, 1 clock = 8ns). (iii). Since our port rate (denoted by R, in Gbps) is constant, we can also directly use it. Based on these information, we can calculate the marking probability given by equation (1). Specifically, when a packet arrives at the switch's output port, the current traffic arriving rate to the output port can be given by P/I. Then the current slope is denoted by $s = P/I - R$. If $0 < s < R$, the marking probability is $Prob = s/R = \frac{P - R \cdot I}{R \cdot I}$. Finally, equation (1) can be rewritten as

$$Prob = \begin{cases} 0, & P \leqslant R \cdot I, \\ \frac{P - R \cdot I}{R \cdot I}, & R \cdot I < P < 2R \cdot I, \\ 1, & P \geqslant 2R \cdot I \end{cases} \quad (2)$$

One method to implement marking with probability is by random number generator. However, using random number generator will introduce large overhead in hardware. Therefore we use another method: instead of marking with probability $Prob$, we can mark every $1/Prob$ packets. Specifically, we can set up a threshold $T = \frac{R \cdot I}{P - R \cdot I}$ and there is a counter counting the number of unmarked packets. When the number of unmarked packets is equal or larger than T, then the next packet will be marked with ECN.

However, calculating the threshold needs to use divider, which will also introduce overhead. In our implementation, we replace divider with adder and comparator. Specifically, let K be the number of unmarked packets, then an arriving packet is marked if $K \geqslant \frac{RI}{P - RI}$. The condition can be rewritten as $K \cdot (P - RI) \geqslant RI$. Instead of counting unmarked packets, we can add up $(P - RI)$ for every arriving packet, and compare the sum with RI. In other words, the arriving packet is marked with ECN if $\sum_{j=1}^{K}(P_j - RI_j) > RI_K$.

Eventually, our implementation structure is shown in Figure 13. When a packet arrives, we calculate $R \cdot I$ and $P - R \cdot I$. Then $\sum_{j=1}^{K}(P_j - RI_j)$ is calculated by add $P - R \cdot I$ to accumulated sum (i.e., sum = sum + $P - RI$). Finally, the ECN marking decision is made by comparing $R \cdot I$ with the sum. Table 1 shows that our implementation needs less than 4% extra resources in NetFPGA.

4.4 Evaluation of S-ECN

In this section, we use realistic experiments in the testbed described in §3.1 to show the performance of the S-ECN scheme.

4.4.1 *Protocols compared*

S-ECN scheme can work with multiple existing end host congestion control algorithms. Specifically, we consider 6 protocols (these protocols are listed in Table 2 for the sake of brevity):

(i) TCP: End hosts use TCP NewReno as their congestion control algorithms. Delayed ACK is disabled [47]. Switch uses taildrop queue and will not mark any packets with ECN.

(ii) ECN*: The protocol proposed in [46]. End hosts use TCP NewReno as their congestion control algorithms, with ECN support (`/proc/sys/net/ipv4/tcp_ecn = 1`). The switch settings are similar as [46]. Specifically, a packet will be marked with ECN when it enters into the output queue and the queue length is larger than ECN threshold. ECN threshold is set to 32KB.

(iii) S-ECN: End hosts settings are the same as those of ECN. Switch use S-ECN scheme to mark packets.

(iv) SL-ECN: S-ECN can suppress the sharp queue length increasing, but may not keep queue length low. Therefore, we introduce SL-ECN protocol, which combines queue-length-based and slope-based ECN marking schemes. Specifically, end host settings are the same

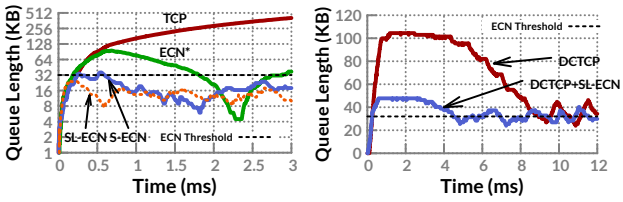

(a) TCP (b) DCTCP

Figure 14: Queue length evolution with difference protocols

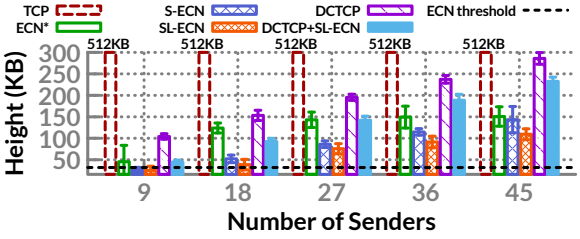

Figure 15: Maximum height the queue length in the congested port can reach. Maximum height of TCP is always 512KB because queue length will not stop increasing until buffer overflows.

as those of ECN*. In the switch, when the queue length is lower than the ECN threshold, switch uses S-ECN scheme to mark packets. When the queue length is higher than the ECN threshold, all packets are marked with ECN. ECN threshold is set to 32KB.

(v) DCTCP: The protocol proposed in [6]. We set $g = 0.125$. Other settings are the same as those of ECN*, including settings of switch.

(vi) DCTCP+SL-ECN: End host settings are the same as those of DCTCP. Switch settings are the same as those of SL-ECN.

4.4.2 *Microbenchmarks*

Sharp queue increasing suppression: First, we show how our scheme can help suppress the sharp queue length increasing. In the experiment, we set the buffer size to 512KB. We let a master host (Host 10) query slave hosts (Host 1-9) for 20MB data and get the instant queue length in the congested port.

The queue length evolution of each protocol is shown in Figure 14. Protocols are divided into two categories by whether end hosts use TCP (Figure 14(a)) or DCTCP (Figure 14(b)) algorithm. First, when end host uses TCP algorithm, there are 4 protocols according to the ECN settings of switch. With TCP, the queue length will be sharply increasing until buffer overflows. With ECN* protocol, the queue length will be sharply increasing even after the queue length exceeds the ECN threshold. The maximum queue length can reach 95.5KB, which is almost three times as high as the ECN threshold (32KB). The reason why the queue length can not be limited under ECN threshold has been analyzed in §4.1. With S-ECN protocol, packets are marked ECN as soon as the queue length begins to increase at the beginning. Therefore the queue length can be quickly controlled. The maximum queue length is only 35.8KB, reduced by 62.5% compared with ECN* protocol. Queue length evolution of SL-ECN protocol is similar as that of SL-ECN protocol, since the queue length for S-ECN protocol is always lower than the ECN threshold, Furthermore, when all flows are in congestion avoidance phase (i.e., after $t = 2$ms), with S-ECN and SL-ECN, the queue length is more steadier than that with ECN*.

The standard deviation of queue length is 4.657KB and 3.383KB with S-ECN and SL-ECN, respectively. While the standard deviation is 13.234KB with ECN protocol. Besides, with S-ECN and SL-ECN, the queue length can keep low. Therefore, the effect of hidden buffer on micro-burst traffic can also be mitigated. When the end hosts use DCTCP algorithm, there are 2 protocols (Figure 14(b)). With DCTCP, the queue length can sharply increase to 104KB[4] at the beginning, although ECN threshold is only 32KB. With DCTCP+SL-ECN the queue length can be controlled under 47KB, reduced by 53.8% compared with DCTCP. Furthermore, when all flows are in congestion avoidance phase (i.e., after $t = 10$ms), with DCTCP+SL-ECN, the queue length is more steadier than that with DCTCP. The standard deviation of queue length is 3.298KB. Compare with it, with DCTCP, the standard deviation is 6.798KB, which is almost twice as large as that with DCTCP+SL-ECN. Besides, as DCTCP+SL-ECN keeps queue length low, the effect of hidden buffer on micro-burst traffic can also be mitigated.

To show the maximum height the queue length can reach when there are more slaves, we enlarge the number of slaves by letting each sending host (i.e., Host 1-9) emulate multiple slaves. We re-conduct the previous experiment when the number of slaves is 9, 18, 27, 36, and 45, and each experiment is repeated for 100 times. Figure 15 shows the maximum height of each protocol. Protocols using S-ECN scheme in the switch can well control the queue length when the number of slaves is large. For example, when there are 36 slaves, compared with ECN*, the maximum queue length can be reduced by 23.5% (S-ECN) and 38.9%(SL-ECN). Compared with DCTCP, the maximum queue length can be reduced by 20.7% with DCTCP+SL-ECN.

Network utilization: In this part, we evaluate the protocols when there are small number of long-lived flows. To show this, we start 3 flows every 1 second,

[4]The maximum queue length is higher than that of ECN* protocol, since DCTCP considers history when judging congestion level and its congestion window decreasing method is smoother.

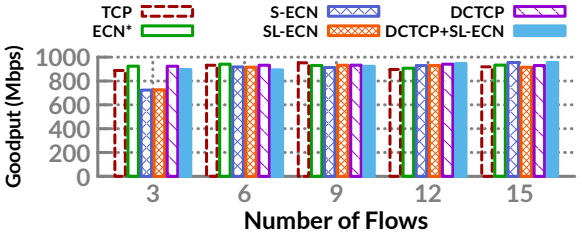

Figure 16: Goodput when the number of flows is small

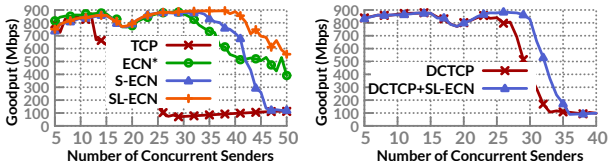

(a) TCP (b) DCTCP

Figure 17: Goodput when each slave respond with 64KB data

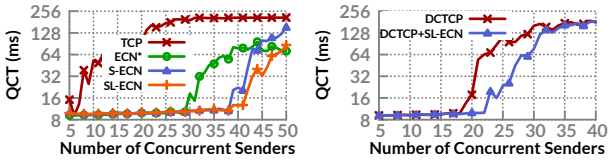

(a) TCP (b) DCTCP

Figure 18: Query Completion Time (QCT) when each slave respond with $1024/n$ KB data. Note that y axis is log-scaled.

who are from different senders and destine for the same receiver. Each flow will send data as fast as it can after it starts.

We obtain the total throughput among all flows when a new bunch of flows start, which is shown in Figure 16. Throughput of DCTCP+SL-ECN is always over 900Mbps. Thus, link utilization is nearly 100%. However, Throughput of S-ECN and SL-ECN is only 723Mbps and 726Mbps, respectively, when there are 3 flows. This is because with these protocols sender will half its congestion window whenever it receives an ECN feedback, which is too much. Since with S-ECN scheme the fraction of marked packets is proportional to slope, it can work better when senders cut its congestion window according to the fraction of marked packets (e.g., DCTCP).

Incast performance: By using S-ECN scheme, flows can avoid massive packet dropping caused by rapid queue length increasing, and thus achieve high throughput. In this part, we'll show how S-ECN scheme improves incast performance.

We consider two cases. The 1st case is common in

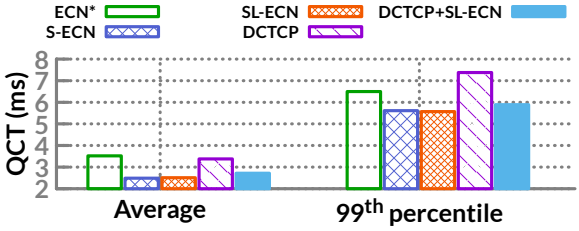

Figure 19: Query completion time.

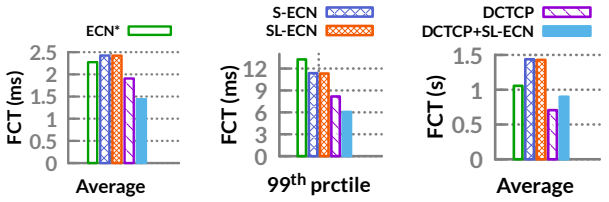

(a) $(0, 100\text{KB}]$ (b) $(0, 100\text{KB}]$ (c) $(10\text{MB}, +\infty)$

Figure 20: Flow completion time across different flow sizes.

MapReduce applications. A master server (Host 10) will send a query to other slaves (Host 1-9). Each slave will respond with 64KB data. Each host is used to emulate multiple senders [44]. The buffer size is 128KB. Figure 17 shows the goodput when the number of slaves varies from 2 to 50. In Figure 17(a), end hosts use TCP algorithm. S-ECN outperforms ECN* when there are less than 40 hosts. However, when the number of hosts is larger than 40, S-ECN is worse than ECN*. This is because S-ECN can control the slope (i.e., queue length increasing rate), but can not decrease the queue length when there are large number of senders. Compared with it, SL-ECN protocol considers both slope and queue length. It can support over 40 senders. While ECN* protocol can support about 30 senders. In Figure 17(b), end hosts use DCTCP algorithm. All protocols can achieve the same goodput when the number of senders are small. Before TCP incast happens, DCTCP protocol can support 25 senders in our experiment, while DCTCP+SL-ECN protocol can support as many as 30 senders.

The 2nd case is common in web search applications The master (Host 10) will query other slaves (Host 1-9) for 1024KB data in all. In other words, each slave will respond with $1024/n$ KB data. Figure 18 shows the query completion time for each protocol, where we have similar observations as previous experiment.

4.4.3 *Benchmark traffic*

Experiment settings: We use a realistic workload to evaluate our scheme, which is derived from data center supporting web search service [6]. In the workload, there are two kinds of traffic between servers. One is

query traffic, where a server will periodically send a query to all other servers and each server will send a response message back after receiving the query message. The total response size is 100KB. Query arrivals follow a Poisson process. The other kind of traffic is background traffic, which follows a one-to-one pattern. The flow size distribution is from [6]. The flow arrivals in each server follow a Poisson process. The ratio of query traffic and background traffic is drawn from [6], and network load in our experiment is 0.4. In our experiments, switch buffer of each port is 128KB, and RTO_{min} is set to 10ms. Each experiment lasts for 5 minutes. Over 350,000 query flows and 140,00 background flows are generated, respectively.

Experiment results: By suppressing the fast queue increasing, flows can finish faster. Figure 19 shows the query completion time in web search workload[5]. Compared to ECN* protocol, S-ECN and SL-ECN can reduce the average query completion time by 29.4% and 28.7%, respectively, and they can reduce the 99^{th} query completion time by 13.6% and 14.3%, respectively. Compared to DCTCP protocol, DCTCP+SL-ECN can reduce the average query completion time by 19.6%, and 99^{th} query completion time is reduced by 20.1%.

Small flows can benefit from suppressing the fast queue increment caused by query traffic. Figure 20 (a,b) shows the average and 99^{th} percentile completion time of flows whose size is smaller than 100KB. Compared to ECN* protocol, S-ECN and SL-ECN do not improve the average flow completion time, but they can reduce the 99^{th} percentile of flow completion time by 14.2% and 14.5%, respectively. Compare to DCTCP protocol, DCTCP+SL-ECN can reduce the average flow completion time by 24.5% and reduce the 99^{th} percentile by 25.9%.

However, S-ECN scheme still needs to be improved. As shown in Figure 20(c), S-ECN, SL-ECN, and DCTCP + SL-ECN protocol do not perform well for large flows. This is part of our future work.

5. RELATED WORK

Lots of studies investigate the burstiness in wide area networks. In [50], authors find that ACK packets can be bunched up when encountering queuing, which in turn causes bursty transmission at sender. The phenomenon is called ACK-compression. The study [11] finds that ACK reordering can result in traffic burstiness. [5] outlines burstiness in TCP and investigates TCP pacing's performance. [26] identifies several causes of bursts from individual flows and examine their effects. [27] examines how TCP creates burstiness in sub-RTT timescales, and finds that TCP's self-clocking and network queuing can shape the packet interarrivals into two-level ON-OFF pattern. [13] assesses the presence and impact of TCP burstiness through analyzing real network traces, and finds that large bursts can always cause packet loss but they rarely occur. Several burst mitigation methods, including MaxBurst, Aggressive Maxburst, User It or Lose It, Congestion Window and Slow Start Threshold Limiting, are studied in [8]. [38] studies the burstiness of TCP flows at packet level. However, these studies mainly focus on burstiness in a single flow. While we study the behavior of micro-burst traffic when many flows aggregate in data center networks.

In data centers, traces from ten data centers is examined in [12], and authors find that traffic exhibits an ON/OFF pattern. [30] studies new causes of bursts, including offloading features in NIC, batching schemes, and bursty OS system calls and APIs. Several tools [40, 25] are developed to detect micro-bursts. In [23], authors find that short-lived congestion caused by micro-burst brings challenges to load balancing systems. Different from these studies, we focus on the nature of micro-burst traffic introduced by transport protocol. When micro-burst traffic occurs, lots of solutions are proposed to address the TCP incast problem [42, 45, 49, 10], and reduce flow completion time [6, 44, 41, 48, 15]. However, these solutions do not touch the essence of the micro-burst traffic.

6. CONCLUSION

Micro-burst traffic in data centers may cause packet dropping, which can bring about serious performance degradation (e.g., TCP incast problem). However, the root cause and dynamic behavior of micro-burst are not comprehensively studied under data center network's unique communication patterns and topologies. In this paper, we use real experiments to examine the micro-burst traffic through observing queue length at find-grained timescale in typical traffic scenarios. We find that self-clocking mechanism and bottleneck link jointly dominate the evolution of micro-burst. We also find that slope of queue length increasing can describe the dynamic behavior of micro-burst under all scenarios. These findings implicates that traditional solutions like absorbing and pacing to mitigate micro-burst traffic are not effective in data centers. To suppress the micro-burst traffic, the sending rate need to be throttled as soon as possible. Enlightened by our findings and implications, we propose S-ECN policy, which takes advantage of slope of queue length increasing and randomly marks packets with ECN. The experiment results show that protocols using S-ECN policy can effectively miti-

[5]Since the performance of TCP is outside the plot range, it is omitted in the figure.

gate micro-burst traffic.

7. REFERENCES

[1] A TCP pacing implementation for NS2. http://netlab.caltech.edu/projects/ns2tcplinux/ns2pacing/.

[2] ns-2. http://www.isi.edu/nsnam/ns/.

[3] Performance Management for Latency-Intolerant Financial Trading Networks. Financial Service Technology.

[4] The NetFPGA Project. http://netfpga.org/.

[5] A. Aggarwal, S. Savage, and T. Anderson. Understanding the performance of TCP pacing. In *Proc. IEEE INFOCOM*, volume 3, pages 1157–1165 vol.3, Mar 2000.

[6] M. Alizadeh, A. Greenberg, D. A. Maltz, J. Padhye, P. Patel, B. Prabhakar, S. Sengupta, and M. Sridharan. Data Center TCP (DCTCP). In *Proceedings of the ACM SIGCOMM 2010 Conference*, SIGCOMM '10, pages 63–74. ACM, 2010.

[7] M. Alizadeh, A. Kabbani, T. Edsall, B. Prabhakar, A. Vahdat, and M. Yasuda. Less is More: Trading a Little Bandwidth for Ultra-low Latency in the Data Center. In *Proc. USENIX NSDI*, pages 19–19, Berkeley, CA, USA, 2012.

[8] M. Allman and E. Blanton. Notes on Burst Mitigation for Transport Protocols. *SIGCOMM Comput. Commun. Rev.*, 35(2):53–60, Apr. 2005.

[9] Arista Networks. Microbursts, Jitter and Buffers. White Paper, 2014.

[10] W. Bai, K. Chen, H. Wu, W. Lan, and Y. Zhao. PAC: Taming TCP Incast Congestion Using Proactive ACK Control. In *Proc. ICNP*, pages 385–396, Oct 2014.

[11] J. Bennett, C. Partridge, and N. Shectman. Packet reordering is not pathological network behavior. *Networking, IEEE/ACM Transactions on*, 7(6):789–798, Dec 1999.

[12] T. Benson, A. Akella, and D. A. Maltz. Network Traffic Characteristics of Data Centers in the Wild. In *Proc. ACM IMC*, pages 267–280. ACM, 2010.

[13] E. Blanton and M. Allman. On the Impact of Bursting on TCP Performance. In C. Dovrolis, editor, *PAM*, volume 3431, pages 1–12. Springer Berlin Heidelberg, 2005.

[14] Y. Chen, R. Griffith, J. Liu, R. H. Katz, and A. D. Joseph. Understanding TCP Incast Throughput Collapse in Datacenter Networks. In *Proc. ACM WREN*, pages 73–82. ACM, 2009.

[15] P. Cheng, F. Ren, R. Shu, and C. Lin. Catch the Whole Lot in an Action: Rapid Precise Packet Loss Notification in Data Centers. In *Proc. USENIX NSDI*, pages 17–28, Berkeley, CA, USA, 2014.

[16] K. Christensen, P. Reviriego, B. Nordman, M. Bennett, M. Mostowfi, and J. Maestro. IEEE 802.3az: the road to energy efficient ethernet. *Communications Magazine, IEEE*, 48(11):50–56, November 2010.

[17] W. chun Feng, P. Tinnakornsrisuphap, and I. Philip. On the burstiness of the tcp congestion-control mechanism in a distributed computing system. In *Distributed Computing Systems, 2000. Proceedings. 20th International Conference on*, pages 110–117, 2000.

[18] Cisco Systems, Inc. Monitor Microbursts on Cisco Nexus 5600 Platform and Cisco Nexus 6000 Series Switches. White Paper, October 2014.

[19] S. Das and R. Sankar. Broadcom smart-buffer technology in data center switches for cost-effective performance scaling of cloud applications. *Broadcom White Paper*, 2012.

[20] J. Dean and S. Ghemawat. MapReduce: Simplified Data Processing on Large Clusters. *Commun. ACM*, 51(1):107–113, Jan. 2008.

[21] Extreme Networks. Congestion Management and Buffering in Data Center Networks. White Paper, 2014.

[22] S. Floyd and V. Jacobson. Random early detection gateways for congestion avoidance. *Networking, IEEE/ACM Transactions on*, 1(4):397–413, Aug 1993.

[23] S. Ghorbani, B. Godfrey, Y. Ganjali, and A. Firoozshahian. Micro Load Balancing in Data Centers with DRILL. In *Proceedings of the 14th ACM Workshop on Hot Topics in Networks*, HotNets-XIV, pages 17:1–17:7, New York, NY, USA, 2015. ACM.

[24] C. Hollot, V. Misra, D. Towsley, and W.-B. Gong. On designing improved controllers for AQM routers supporting TCP flows. In *Proc. IEEE INFOCOM*, volume 3, pages 1726–1734 vol.3, 2001.

[25] V. Jeyakumar, M. Alizadeh, Y. Geng, C. Kim, and D. Mazières. Millions of Little Minions: Using Packets for Low Latency Network Programming and Visibility. In *Proceedings of the 2014 ACM Conference on SIGCOMM*, SIGCOMM '14, pages 3–14. ACM, 2014.

[26] H. Jiang and C. Dovrolis. Source-level IP Packet Bursts: Causes and Effects. In *Proc. ACM IMC*, pages 301–306. ACM, 2003.

[27] H. Jiang and C. Dovrolis. Why is the Internet Traffic Bursty in Short Time Scales? In *Proc. ACM SIGMETRICS*, pages 241–252, New York, NY, USA, 2005. ACM.

[28] S. Kandula, S. Sengupta, A. Greenberg, P. Patel, and R. Chaiken. The Nature of Data Center

Traffic: Measurements & Analysis. In *Proc. ACM IMC*, pages 202–208. ACM, 2009.

[29] R. Kapoor, G. Porter, M. Tewari, G. M. Voelker, and A. Vahdat. Chronos: Predictable Low Latency for Data Center Applications. In *Proceedings of the Third ACM Symposium on Cloud Computing*, SoCC '12, pages 9:1–9:14, New York, NY, USA, 2012. ACM.

[30] R. Kapoor, A. C. Snoeren, G. M. Voelker, and G. Porter. Bullet Trains: A Study of NIC Burst Behavior at Microsecond Timescales. In *Proceedings of the Ninth ACM Conference on Emerging Networking Experiments and Technologies*, CoNEXT '13, pages 133–138, New York, NY, USA, 2013. ACM.

[31] L. Mai, C. Hong, and P. Costa. Optimizing network performance in distributed machine learning. In *Proceedings of the 7th USENIX Conference on Hot Topics in Cloud Computing*, pages 2–2. USENIX Association, 2015.

[32] S. Makineni, R. Iyer, P. Sarangam, D. Newell, L. Zhao, R. Illikkal, and J. Moses. Receive Side Coalescing for Accelerating TCP/IP Processing. In Y. Robert, M. Parashar, R. Badrinath, and V. Prasanna, editors, *High Performance Computing - HiPC 2006*, volume 4297 of *Lecture Notes in Computer Science*, pages 289–300. Springer Berlin Heidelberg, 2006.

[33] L. Neumeyer, B. Robbins, A. Nair, and A. Kesari. S4: Distributed Stream Computing Platform. In *Data Mining Workshops (ICDMW), 2010 IEEE International Conference on*, pages 170–177, Dec 2010.

[34] R. Nishtala, H. Fugal, S. Grimm, M. Kwiatkowski, H. Lee, H. C. Li, R. McElroy, M. Paleczny, D. Peek, P. Saab, D. Stafford, T. Tung, and V. Venkataramani. Scaling Memcache at Facebook. In *Proc. USENIX NSDI*, pages 385–398, Berkeley, CA, USA, 2013.

[35] A. Phanishayee, E. Krevat, V. Vasudevan, D. G. Andersen, G. R. Ganger, G. A. Gibson, and S. Seshan. Measurement and Analysis of TCP Throughput Collapse in Cluster-based Storage Systems. In *FAST*, volume 8, pages 1–14, 2008.

[36] K. Ramakrishnan, S. Floyd, D. Black, et al. The Addition of Explicit Congestion Notification (ECN) to IP. RFC 3168, RFC Editor, September 2001.

[37] J. Rothschild. High performance at massive scale: Lessons learned at facebook. https://www.youtube.com/watch?v=KIRzk08NMNo, 2009.

[38] S. Shakkottai, N. Brownlee, and k. claffy. A Study of Burstiness in TCP Flows. In C. Dovrolis, editor, *PAM*, volume 3431 of *Lecture Notes in Computer Science*, pages 13–26. Springer Berlin Heidelberg, 2005.

[39] D. Shan, W. Jiang, and F. Ren. Absorbing micro-burst traffic by enhancing dynamic threshold policy of data center switches. In *Proc. IEEE INFOCOM*, pages 118–126, April 2015.

[40] F. Uyeda, L. Foschini, F. Baker, S. Suri, and G. Varghese. Efficiently Measuring Bandwidth at All Time Scales. In *Proc. USENIX NSDI*, pages 71–84, Berkeley, CA, USA, 2011.

[41] B. Vamanan, J. Hasan, and T. Vijaykumar. Deadline-aware Datacenter TCP (D2TCP). In *Proc. of ACM SIGCOMM*, pages 115–126. ACM, aug 2012.

[42] V. Vasudevan, A. Phanishayee, H. Shah, E. Krevat, D. G. Andersen, G. R. Ganger, G. A. Gibson, and B. Mueller. Safe and Effective Fine-grained TCP Retransmissions for Datacenter Communication. In *Proc. ACM SIGCOMM*, pages 303–314. ACM, 2009.

[43] D. Wei, P. Cao, S. Low, and C. EAS. TCP pacing revisited. unpublished.

[44] C. Wilson, H. Ballani, T. Karagiannis, and A. Rowtron. Better Never Than Late: Meeting Deadlines in Datacenter Networks. In *Proc. of ACM SIGCOMM*, pages 50–61. ACM, 2011.

[45] H. Wu, Z. Feng, C. Guo, and Y. Zhang. ICTCP: Incast Congestion Control for TCP in Data-center Networks. *IEEE/ACM Trans. Netw.*, 21(2):345–358, Apr. 2013.

[46] H. Wu, J. Ju, G. Lu, C. Guo, Y. Xiong, and Y. Zhang. Tuning ECN for Data Center Networks. In *Proceedings of the 8th International Conference on Emerging Networking Experiments and Technologies*, CoNEXT '12, pages 25–36, New York, NY, USA, 2012. ACM.

[47] M. Yu, A. Greenberg, D. Maltz, J. Rexford, L. Yuan, S. Kandula, and C. Kim. Profiling Network Performance for Multi-tier Data Center Applications. In *Proc. USENIX NSDI*, pages 57–70, Berkeley, CA, USA, 2011.

[48] D. Zats, T. Das, P. Mohan, D. Borthakur, and R. Katz. DeTail: Reducing the Flow Completion Time Tail in Datacenter Networks. In *Proc. of ACM SIGCOMM*, pages 139–150. ACM, 2012.

[49] J. Zhang, F. Ren, L. Tang, and C. Lin. Taming TCP incast throughput collapse in data center networks. In *Proc. ICNP*, pages 1–10, Oct 2013.

[50] L. Zhang, S. Shenker, and D. D. Clark. Observations on the Dynamics of a Congestion Control Algorithm: The Effects of Two-way Traffic. In *Proc. ACM SIGCOMM*, pages 133–147. ACM, 1991.